

LSI 01-674

PATENT

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

APPLICATION FOR PATENT

ON

DUAL PORTING SERIAL ADVANCED TECHNOLOGY ATTACHMENT DISK
DRIVES FOR FAULT TOLERANT APPLICATIONS

BY

BRET S. WEBER
2521 N. TEE TIME
WICHITA, KS 67205
CITIZEN OF USA

JOHN V. SHERMAN
1696 TIARA PINES CRC
DERBY, KS 67037
CITIZEN OF USA

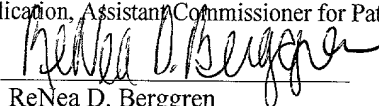
CERTIFICATE OF MAILING BY "EXPRESS MAIL"

"Express Mail" Mailing Label Number EV 013 245 069 US

Date of Deposit: December 19, 2001

I hereby certify that this correspondence is being deposited with the United States Postal Service "Express Mail Post Office to Addressee" service under 37 C.F.R. § 1.10 on the date indicated above and is addressed to Box Patent Application, Assistant Commissioner for Patents, Washington, D.C. 20231

BY:


ReNea D. Berggren

DUAL PORTING SERIAL ATA DISK DRIVES FOR FAULT TOLERANT APPLICATIONS

FIELD OF THE INVENTION

[0001] The present invention generally relates to the field of data storage, and particularly to a system for utilizing serial disk drives in fault tolerant multi-initiated applications.

BACKGROUND OF THE INVENTION

[0002] Data storage has become increasingly important as information handling systems become capable of handling and communicating more data. Currently, improvements to data storage systems are aimed at increasing the amount of data and providing fault tolerant systems. While data storage systems have become more efficient at handling large amounts of data in a fault tolerant manner, typically such fault-tolerant data storage systems often are cost-prohibitive.

[0003] Expensive fault tolerant data storage systems therefore may not be implemented in certain situations where their cost is not justified by the design application. As a result, information handling systems often are configured with lower cost data storage systems. One example of this cost savings measure is in the implementation of serial advanced technology attachment (SATA) disk drives in place of fibre channel disk drives.

[0004] Current fibre channel disk drives are multi-ported and offer fault tolerant interfacing with external components. One of the drawbacks of including a fibre channel disk drive in a system is the drive's price difference. Presently fibre channel disk drives typically are between 3 to 6 times the price of a SATA disk drive, for a comparable storage capacity. Thus designers often may include a SATA drive instead of a fibre channel drive in certain applications.

[0005] While serial disk drive offer increased cost savings over fibre channel drives, the inherent single ported design are often a source of failure, thus, reducing the availability of data. It is this single point of failure which may cause designers either to include costly fibre channel disk drives or risk potential data loss.

[0006] Therefore, it would be desirable to provide a system and method for cost-effectively utilizing serial storage devices in a fault tolerant applications, thereby allowing for the utilization of more cost effective serial storage devices without the drawbacks inherent to a single ported system.

SUMMARY OF THE INVENTION

[0007] Accordingly, the present invention is directed to an apparatus and method for implementing dual porting of an inherently single ported serial disk drive, such as a SATA disk drive. The inclusion of an apparatus of the present invention allows a single ported serial disk drive to operate, in a fault-tolerant system such as a system employing fibre channel based communication without the previously experienced fault drawbacks.

[0008] An apparatus for dual porting a serial disk drive of the present invention in the presently preferred embodiment is connected through the back panel of an information handling system utilizing a peripheral component interface (PCI) to a first and a second SATA master control devices which are operating in a fibre channel based communication system. In the present embodiment the dual porting apparatus is included in a disk drive canister board additionally containing a SATA disk drive. Further by including the dual porting apparatus in the disk drive canister board in the event of a failure the dual porting apparatus may be easily replaced.

[0009] In a further embodiment the dual porting apparatus contains a first and a second idle regenerator connected to a first and a second SATA master device, respectively. The first and the second idle regenerators being capable of transmitting and receiving signals to their respective SATA master devices. Included in the dual porting apparatus is a third idle regenerator connected to the first and the second idle regenerators. The third idle regenerator additionally connected to a serial disk drive, such as a SATA disk drive. The third idle regenerator is capable of receiving and transmitting signals to the serial disk drive. Further included in the dual porting apparatus is synchronization logic capable of synchronizing the transfer between one of the first idle regenerator and the second idle regenerator and the third idle regenerator.

[0010] Additionally, the dual porting apparatus, in contemplated embodiments, is capable of receiving control signals, thereby controlling which idle generator is enabled from a SATA master device. In further embodiments the apparatus for dual porting a serial disk drive includes an auto detector capable of determining the presence/absence of idle characters from the first and the second SATA master devices. Thus the auto detector may determine which SATA master device is active and thus enable a single SATA master device at a time.

[0011] It is to be understood that both the forgoing general description and the following detailed description are exemplary and explanatory only and are not restrictive of the invention as claimed. The accompanying drawings, which are incorporated in and constitute a part of the specification, illustrate an embodiment of the invention and together with the general description, serve to explain the principles of the invention.

BRIEF DESCRIPTION OF THE DRAWINGS

[0012] The numerous advantages of the present invention may be better understood by those skilled in the art by reference to the accompanying figures in which:

[0013] FIG. 1 is an overview illustration of an exemplary embodiment wherein an apparatus for dual porting a serial disk drive in an information handling system;

[0014] FIG. 2 is an overview illustration of an exemplary embodiment an apparatus for dual porting a serial disk drive in a disk drive canister board;

[0015] FIG. 3 is an overview illustration of an apparatus for dual porting a serial disk drive with capability for accepting control signal from an external source.

DETAILED DESCRIPTION OF THE INVENTION

[0016] Reference will now be made in detail to the presently preferred embodiments of the present invention, examples of which are illustrated in the accompanying drawings.

[0017] Referring generally now to FIGS. 1 through 3, exemplary embodiments of the present invention are discussed, wherein an apparatus for utilizing single ported serial disk drives, such as SATA disk drives in multi-ported fault tolerant applications are discussed. Fault tolerant storage systems such as fibre channel disk drives offer increased functionality over SATA drives which are subject to a single point of failure, thus reducing the availability of data to the system.

[0018] While fibre channel drives offer advantages over that of SATA drives, the high cost of fibre channel drives continues to be a factor in overall utilization. For example, fibre channel disk drives are often utilized in high performance information handling systems. Therefore, the current invention provides an apparatus to incorporate SATA drives into fault-tolerant applications without the previously experienced single point failure associated with SATA drives.

[0019] Referring now to FIG.1, in a first embodiment of the present invention an apparatus for dual porting a serial disk drive 102, such as a SATA drive 104, is discussed. In the current embodiment, the dual porting apparatus 102 is included on a disk drive canister 114. The present implementation allows for the utilization of the SATA drive 104 in a fibre channel environment without the drawbacks associated with single point failures. In the present embodiment, the dual porting apparatus 102 is employed as part of the disk drive canister 114 incorporated in an information handling system 100.

[0020] Furthermore, the arrangement of disk drive canister 114 including associated SATA drives 104 and dual porting apparatus 102 allows for redundancy in the information handling system 100, thus reducing the risk of a failure.

[0021] In the present embodiment, the dual porting apparatus 102 is connected to a first and a second SATA master devices 110 & 112 through a back panel 106 of the information handling system 100 through peripheral component interfaces (PCI) 108 implemented in the first and second SATA master devices 110 & 112. It is to be understood that various interfaces may be employed as contemplated by one of ordinary skill in the art without departing from the spirit and scope of the present invention, and it is the intention of the present invention to encompass and include these variations. For example the present interface may be incorporated in a proprietary bus or the like. For present purposes the first and the second SATA master devices are substantially similar and are merely referred to as "first and second" SATA master devices 110 & 112 for convenience.

[0022] In the present embodiment, the first and second SATA master devices 110 & 112 individually include a fibre channel connection 116. Furthermore, the first and second SATA master devices include a uniprocessor (uP) 118 and a memory 120.

[0023] Additionally, in the present embodiment the SATA master devices 110 & 112 are capable of providing control signals from a higher level to the dual porting apparatus 102 ensuring that only a single disk drive canister 114 is enabled at a time. In further embodiments, the first and the second SATA master devices 110 & 112 are continuously enabled with controls executed at a higher level guaranteeing that only a single SATA master device is enabled at a time. For example, by sending control signals from a higher level, switching in the middle of a command is prevented.

[0024] Referring now to FIG. 2, a disk drive canister 200 including a dual porting apparatus 202 for porting a serial disk drive is discussed. In the present embodiment, the dual porting apparatus 202 may be connected between a SATA disk drive 204 and the back panel of an information handling system, such as the arrangement discussed in FIG. 1, thus effectively multiplexing the serial disk drive 204. In the current embodiment the dual porting apparatus 202 is integrated in the disk drive canister 200 thus allowing for ease of replacement in the event of failure.

[0025] The present embodiment thus allows for the incorporation of the SATA disk drive 204 into a fault tolerant environment, such as fibre channel based systems without the drawback of a single failure point previously experienced with SATA drive included without the dual porting apparatus 202 of the present invention. Further, through implementation of the present invention the SATA drive 204 may continue to receive idle characters during utilization.

[0026] The dual porting apparatus 202 includes a first idle regenerator 206. The first idle regenerator 206 is connected to a first serial master device, such as discussed in FIG. 1. The first idle regenerator 206 in the present embodiment is capable of transmitting and receiving signals to the first serial master device via an input/output connection 208. In the current embodiment the input/output connection is made utilizing a PCI, although it is

contemplated that other interfaces are capable of utilization for this purpose without departing from the scope and spirit of the present invention. For example, the interface may be implemented without a processor or be embodied in a proprietary bus, or the like.

[0027] The dual porting apparatus 202 of the present embodiment further includes a second idle regenerator 210. The second idle regenerator 210 is connected to a second serial master device via input/output connection 212. The second serial master device of the present embodiment is substantially similar to the second serial master device as discussed in FIG. 1. The second idle regenerator 210 is capable of receiving and transmitting signals to the second serial master device. Additionally, the second idle regenerator 210 is connected through connection 212 which is ported to the second serial master device through a PCI.

[0028] A third idle regenerator 214 is connected to the first and the second idle regenerators 206 & 210 and additionally to the serial disk drive 204. The third idle regenerator 214, included in the dual porting apparatus 202, is capable of transmitting and receiving signals from the first and second idle regenerators 206 & 210 and the serial disk drive 204, such as a SATA disk drive.

[0029] Further included in the dual porting apparatus 202, is synchronization logic 216 connected to the first, second and third idle regenerators 206, 210 and 216 respectively. The synchronization logic 216 is capable of synchronizing data transfers between one of the first and second idle regenerators 206 & 210 and the third idle generator 214. For example, the synchronization logic 216 is capable of synchronizing the third idle regenerator 214 with the second idle regenerator in the event of a desired switch in data stream input from the first idle regenerator 206 to the second idle regenerator 210. Furthermore, should an event occur which causes a disruption in synchronization in the dual porting apparatus, the synchronization logic 216 is utilized to reestablish

synchronization between the third idle regenerator and the first or second idle regenerator as required by the circumstances.

[0030] In embodiments of the present invention the dual porting apparatus 202 further includes an auto detector 218. The auto detector 218 in the present embodiment is connected to the input/output of the first and the second serial masters devices 208 & 212, respectively as well as the synchronization logic 216.

[0031] For example, the auto detector 218 may be incorporated into the dual porting apparatus 102 which in turn is included in a disk drive canister 114 as discussed in FIG. 1 or the like. The auto detector 218 is capable of automatically switching between the first and second master devices based on the presence or absence of idle characters transmitted by either of the first and second master devices. The auto detector 218 is further capable of controlling the switching between the first and second serial master devices, thus allowing the input/output of data between the dual porting apparatus and one of the first and second master devices at a time.

[0032] In embodiments of the present invention the dual porting apparatus 202 is implemented as an application specific integrated circuit (ASIC).

[0033] Referring generally to FIG.3, in an alternative embodiment a dual porting apparatus 302 may be capable of being switched between the first and the second serial master devices through control signals sent from an external source, such as a SATA master device.

[0034] It is believed that the DUAL PORTING SERIAL ATA DISK DRIVES FOR FAULT TOLERANT APPLICATIONS of the present invention and many of its attendant advantages will be understood by the forgoing description. It is also believed that it will be apparent

that various changes may be made in the form, construction and arrangement of the components thereof without departing from the scope and spirit of the invention or without sacrificing all of its material advantages. The form herein before described being merely an explanatory embodiment thereof. It is the intention of the following claims to encompass and include such changes.